Priorities in Energy Optimized Processing

John Cornish

VP Marketing, Processor Division, ARM

11 October, 2006

Constrained Power Footprint

- Everything has improved except the battery!
  - My phone in 1998: Nokia 6110 270h standby 900mAh NiMH
  - My current phone: iMate SP5 250h standby 1150mAh Lithium-ion

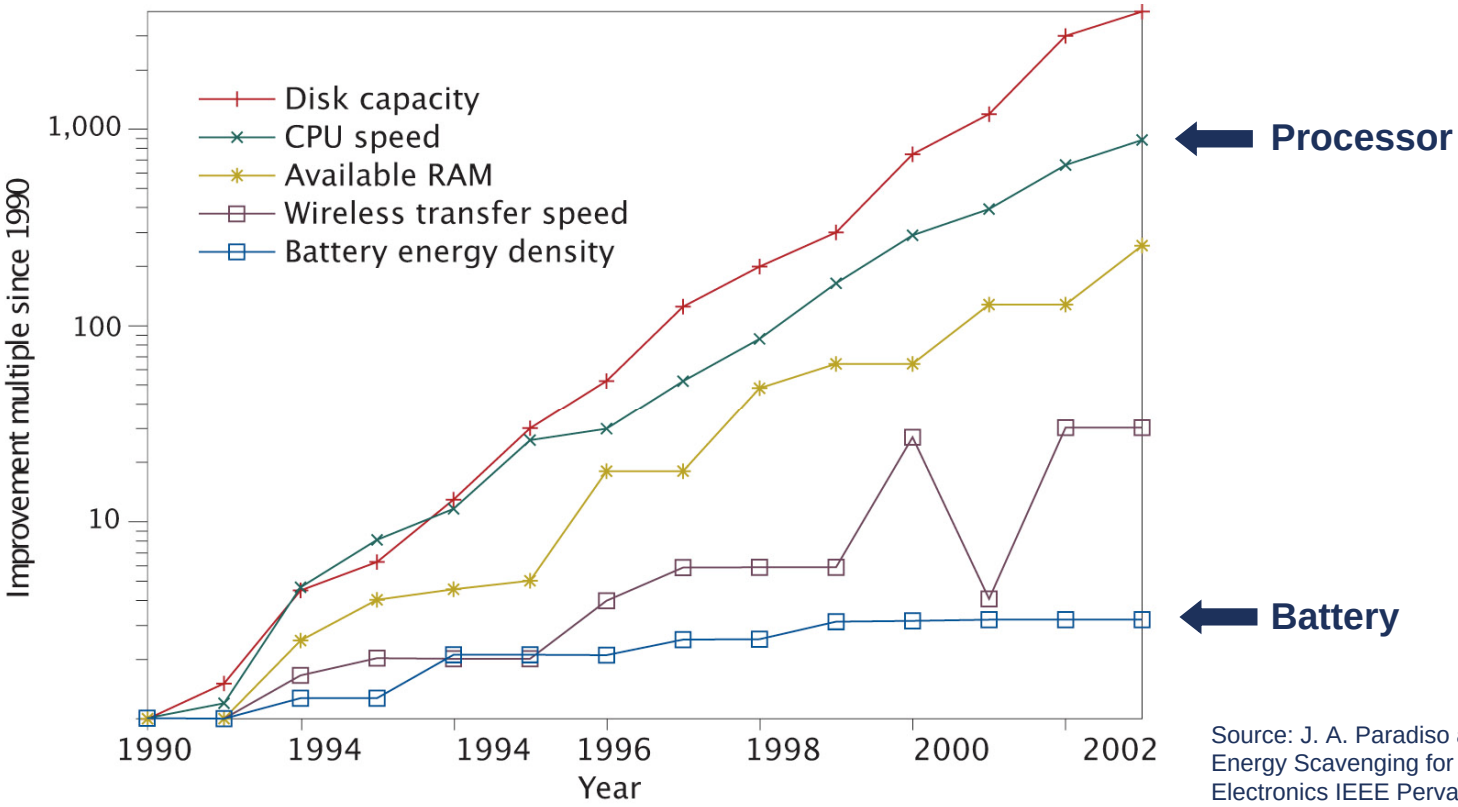

Processor

Battery

Source: J. A. Paradiso and T. Starner, Energy Scavenging for Mobile and Wireless, Electronics IEEE Pervasive Computing, 2005.

Exploiting Process Technology

Cortex-A8
2000 DMIPS
<300mW

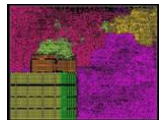

Uncertainty
Variability

ARM11
920 DMIPS
<350mW

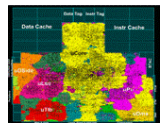

Leakage Power
Lithography

ARM10E
360 DMIPS
<160mW

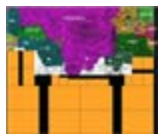

Signal Integrity
CMP

ARM9E
220 DMIPS
<225mW

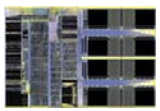

Timing Closure
Vias

ARM7
60 DMIPS
<50mW

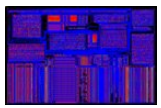

Timing, Area
Particles

250nm 180nm 130nm 90nm 65nm 45nm

Complexity and Power Explosion

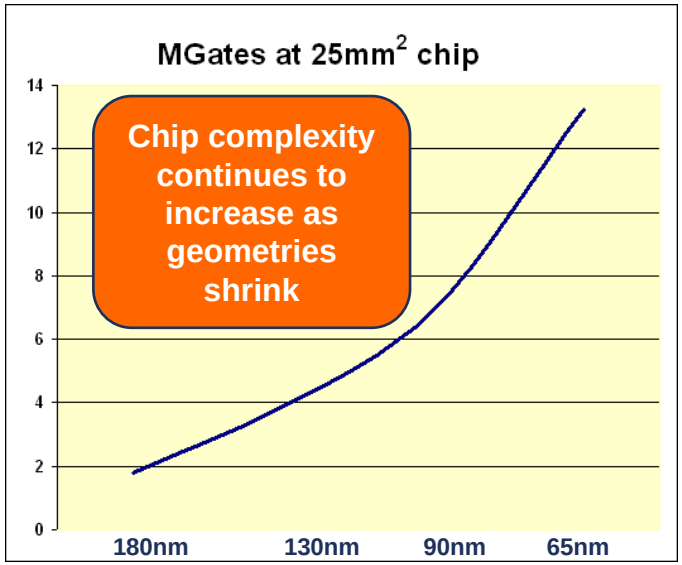

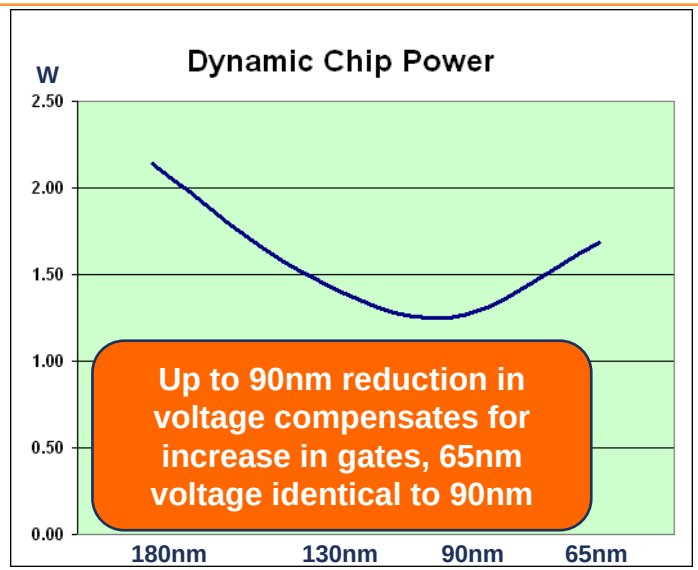

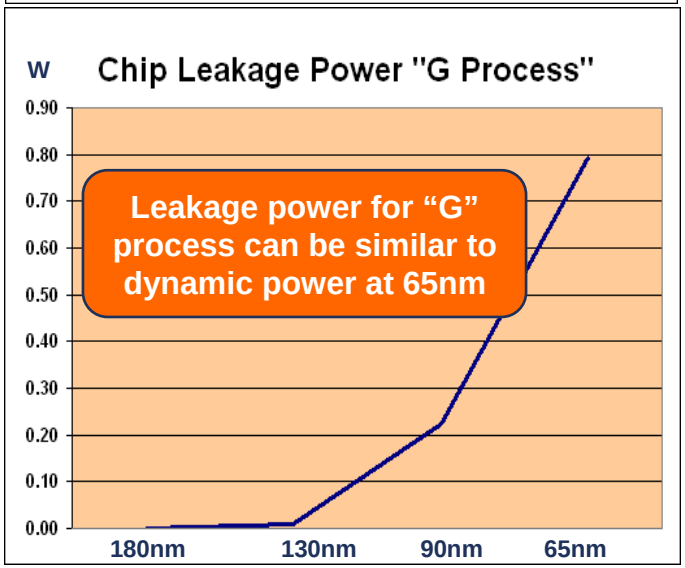

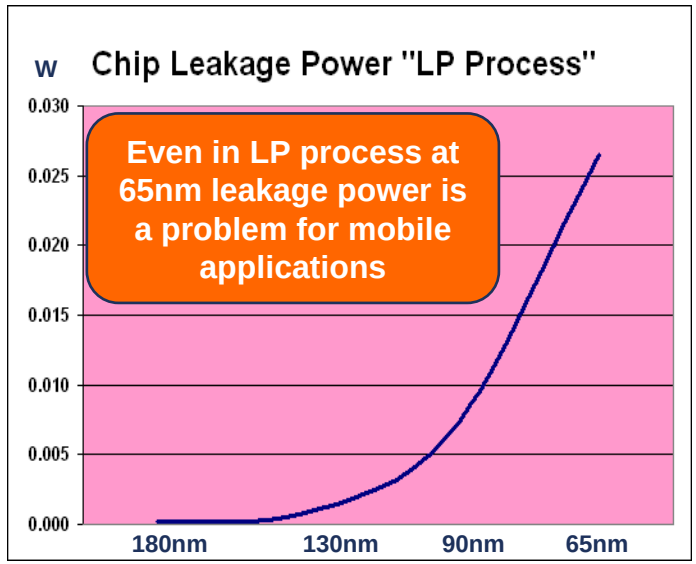

1. Gates per chip based on 75% standard cells and 80% utilization for 25mm² chip per process node,
2. Power based on number of gates from 1, dynamic power with 250MHz and 30% activity, leakage power at 125C, FF, VDD+10%

Multiprocessing (MP)

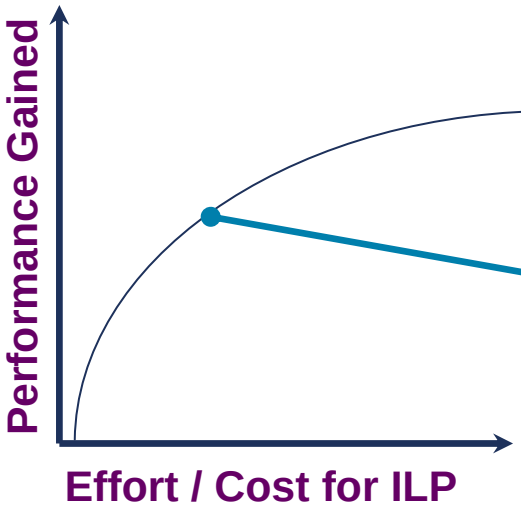

‘Optimal’ support for both ILP and TLP brings the most performance for the least cost / effort

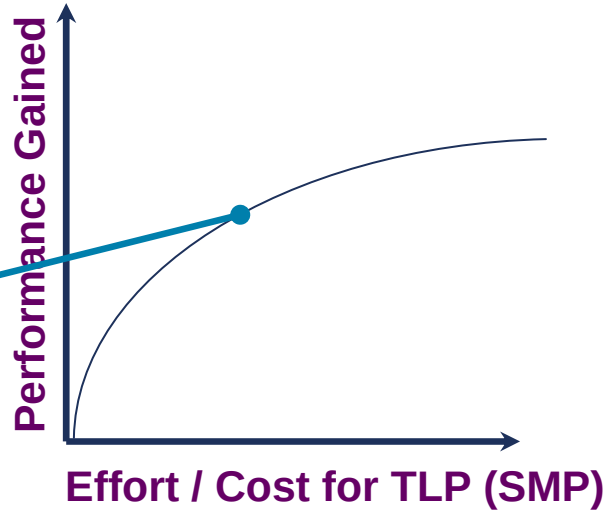

Expensive on hardware to continue to attempt to extract instruction level parallelism (ILP)

Expensive on hardware and software to continue to attempt to share work between more instances of thread level parallelism (TLP)

MP – Exploiting Concurrency

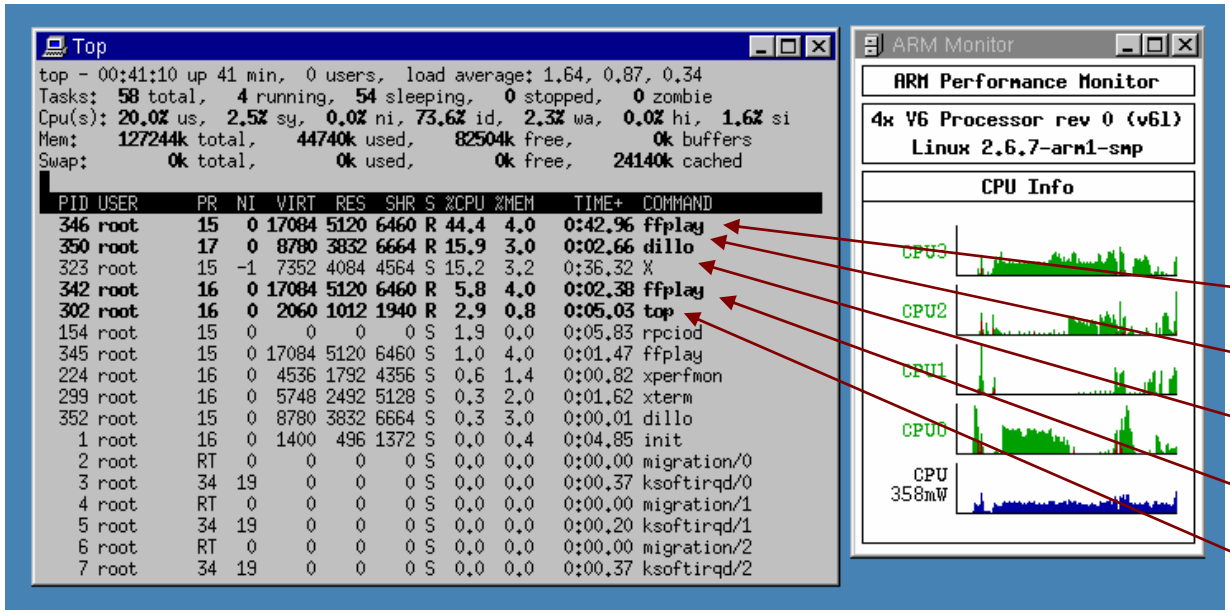

- Multi-tasking
  - Inherent within applications and the operating system
- Multi-threading
  - Offloading tasks to specialist processors
  - Creating a ‘pool’ of tasks the OS can share across general SMP processor
  - Exposing utility tasks that can be scheduled in the background

Power Gating

- Coarse grain power switches
  - Switching power or ground for multiple cells at block level
  - Disconnect local power from global power via power gates (MTCMOS switches)
- Header and Footer switches allow flexibility
  - Header to switch VDD, Footer to switch VSS
  - Same cell height as normal cells
  - Different transistor sizes (cell width) available

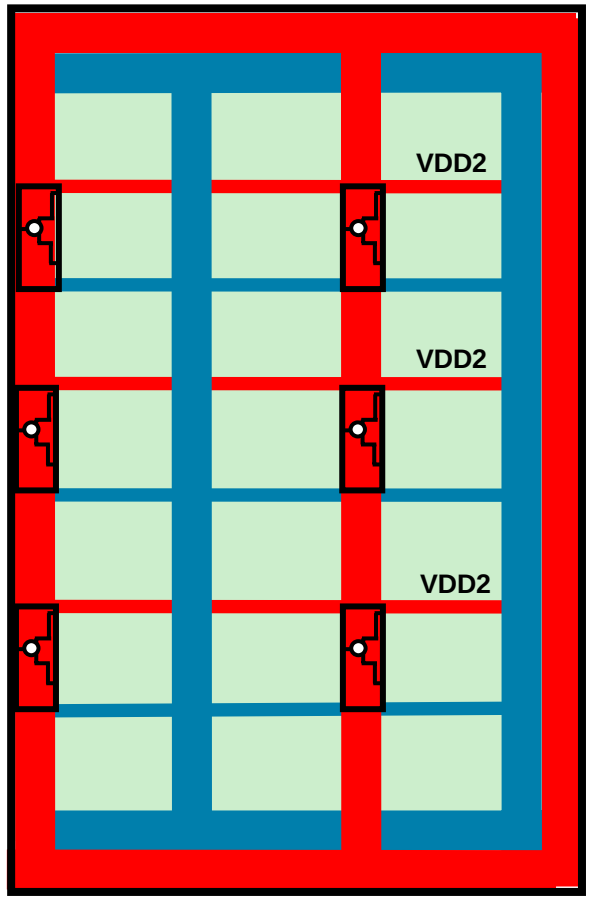

Switches on the connection power strip and power rail to switch VDD2

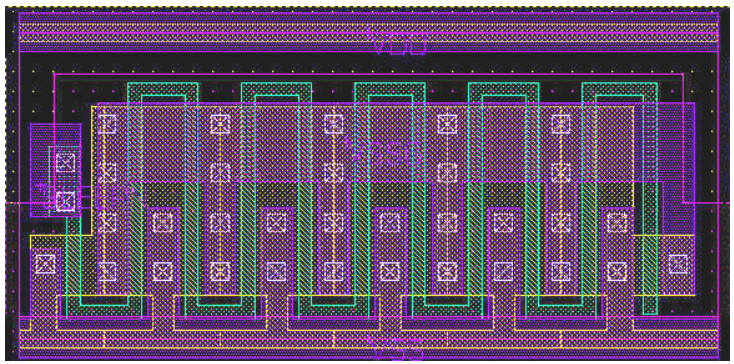

Example Footer Switch

State Retention

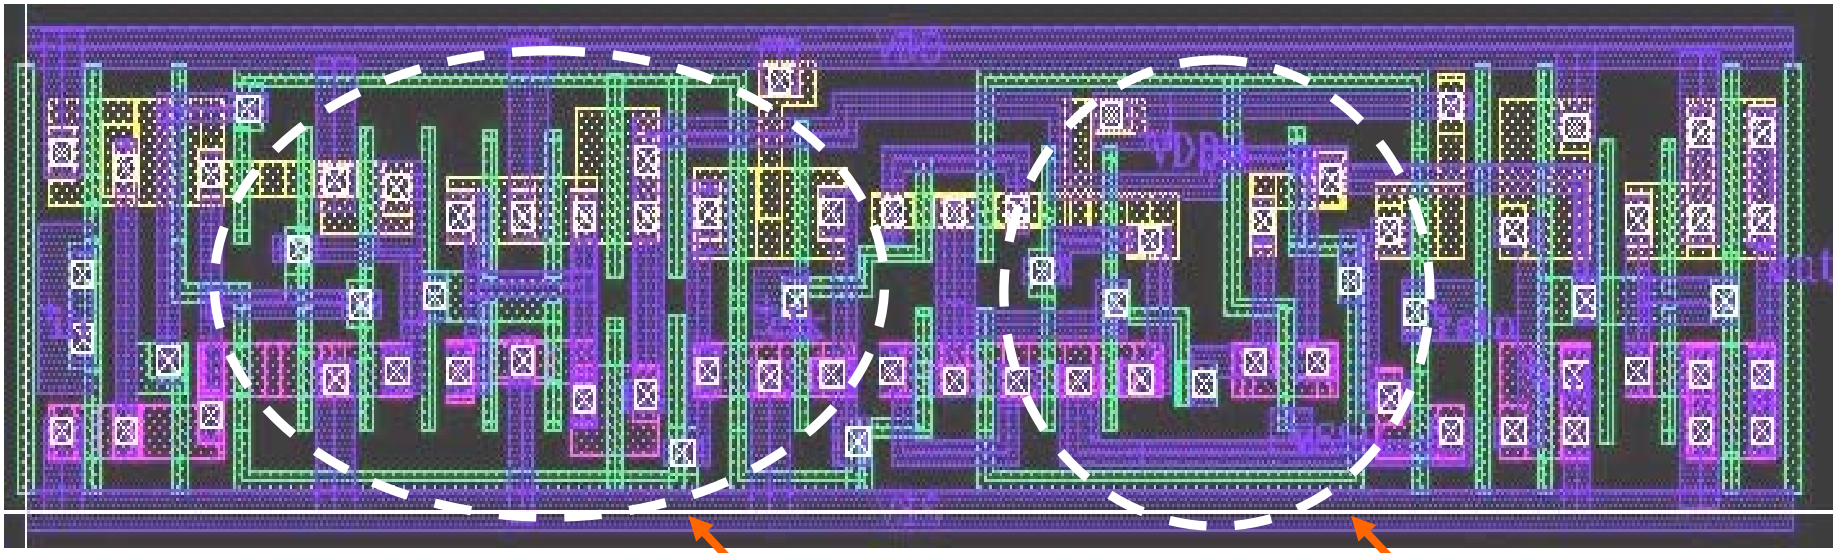

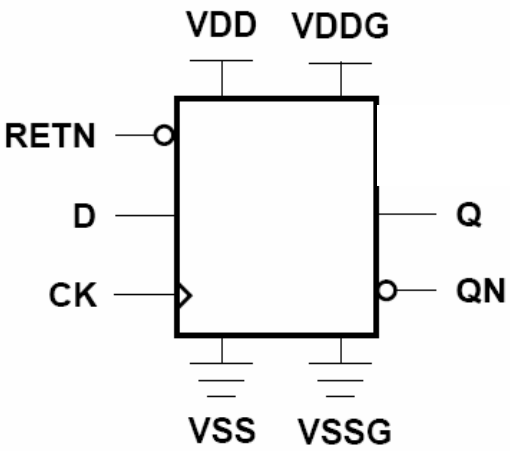

1. D-flop main stage (SVt or LVt)

- No leakage when retaining state
- Connects to switchable (local) power rails

2. D-flop retention stage (HVt)

- Low leakage when retaining state
- Connects to global (always-on) power

65LP Leakage Analysis

- Correlating measurements
  - IV analysis for CPU+RAM
    RAM requires 1.08V min
- HALT
  - Std cell + RAM leakage
- LIGHT SLEEP
  - Lvt Std cell power gated
  - Retention registers active
  - + RAM leakage
- DEEP SLEEP
  - Scan-based hibernate
  - Just RAM leakage

All at room temp (23C)

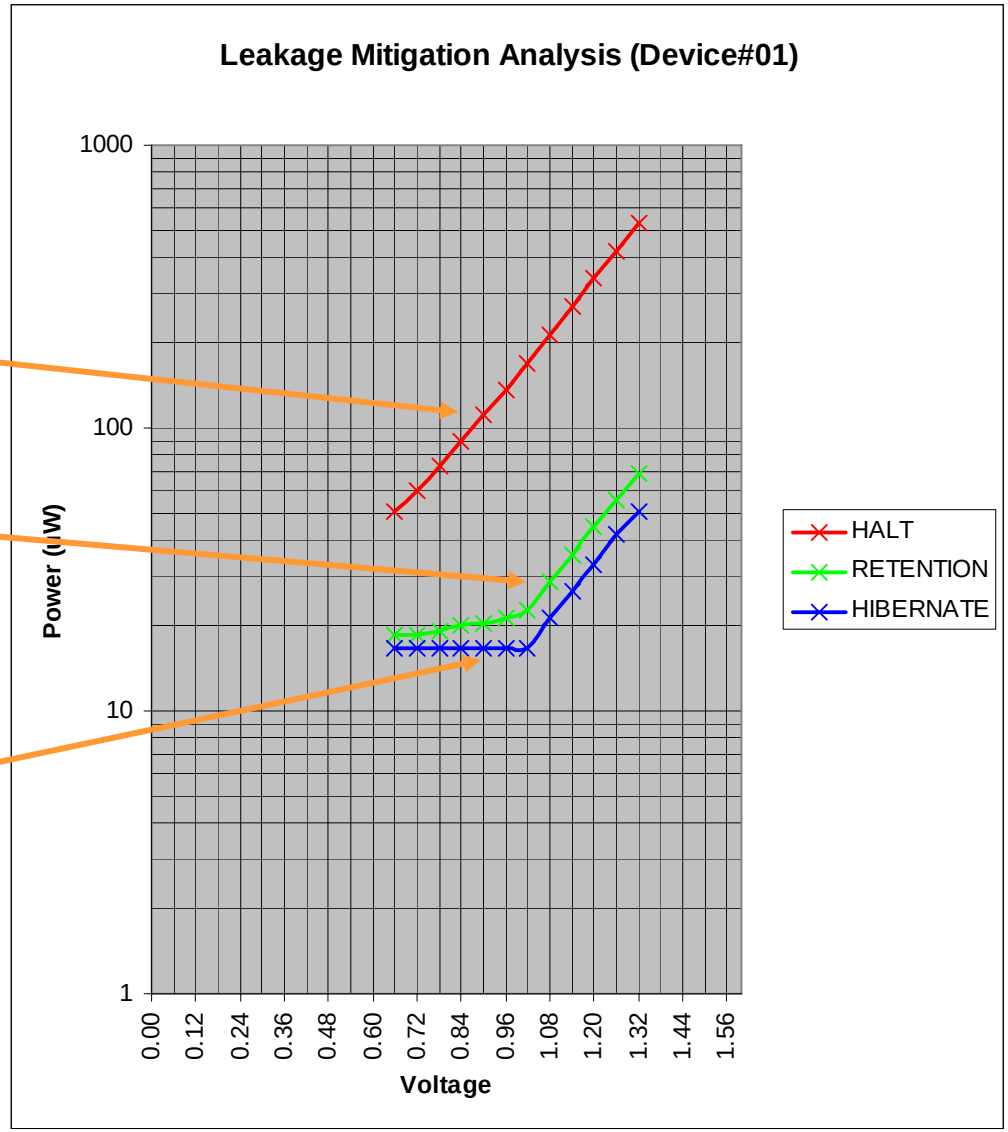

MP – Exploiting Power Gating

- Enter Standby using wait-for-interrupt
  - Maintain cache coherence while in standby
  - Allows immediate re-entry without any cache house keeping
- Enter Dormant/Shutdown state
  - Remove supply to core logic but maintain RAM
  - Remove supply and save dynamic and leakage power
  - Requires all dirty cache data to be evicted prior to entry
  - Awoken CPU reloads state and rejoins the coherence domain

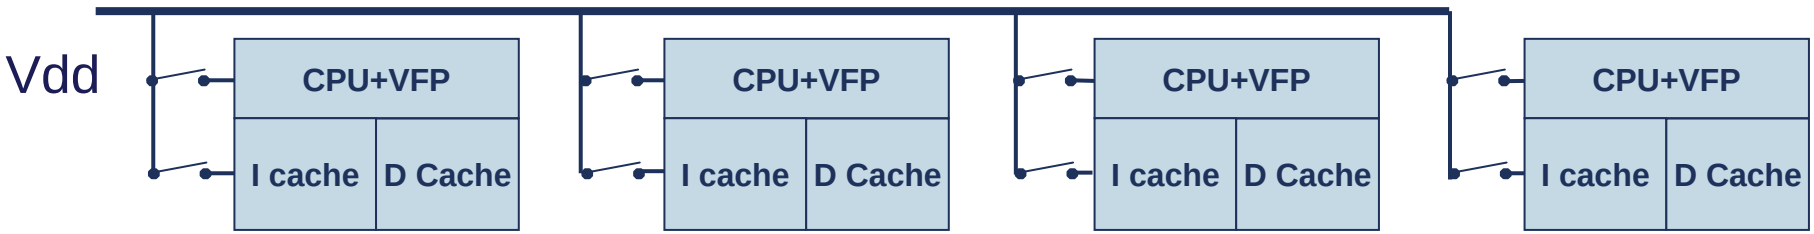

Dormant – Core logic powered down, cache retains contents
Shutdown – Core logic ***and*** cache powered down

Dynamic Voltage & Frequency Scaling

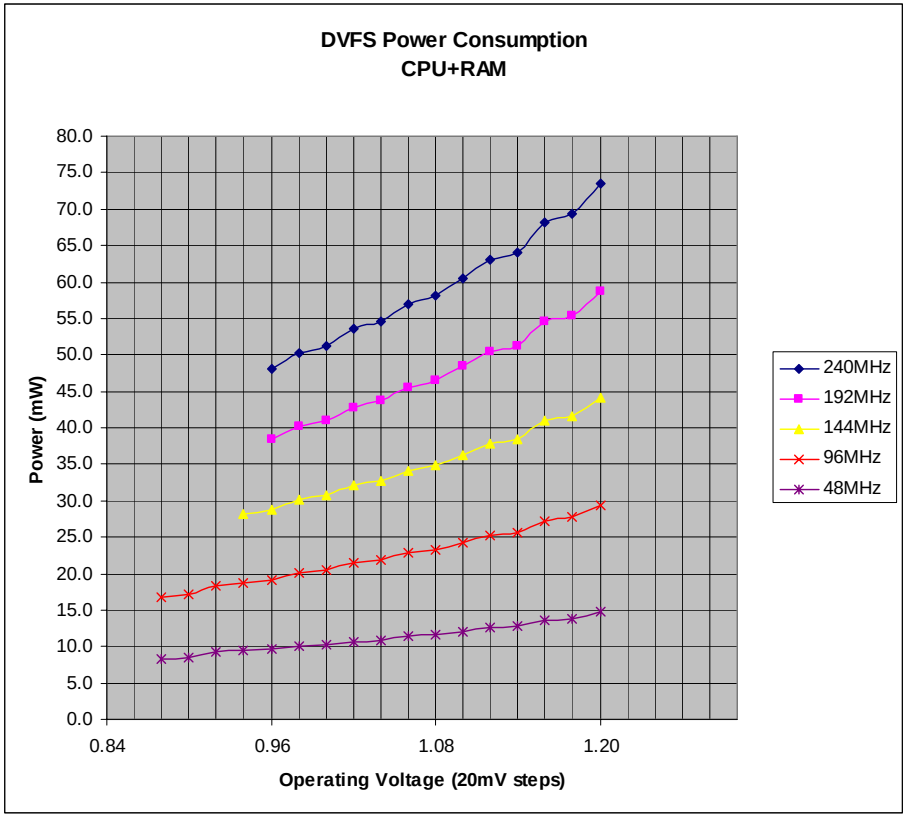

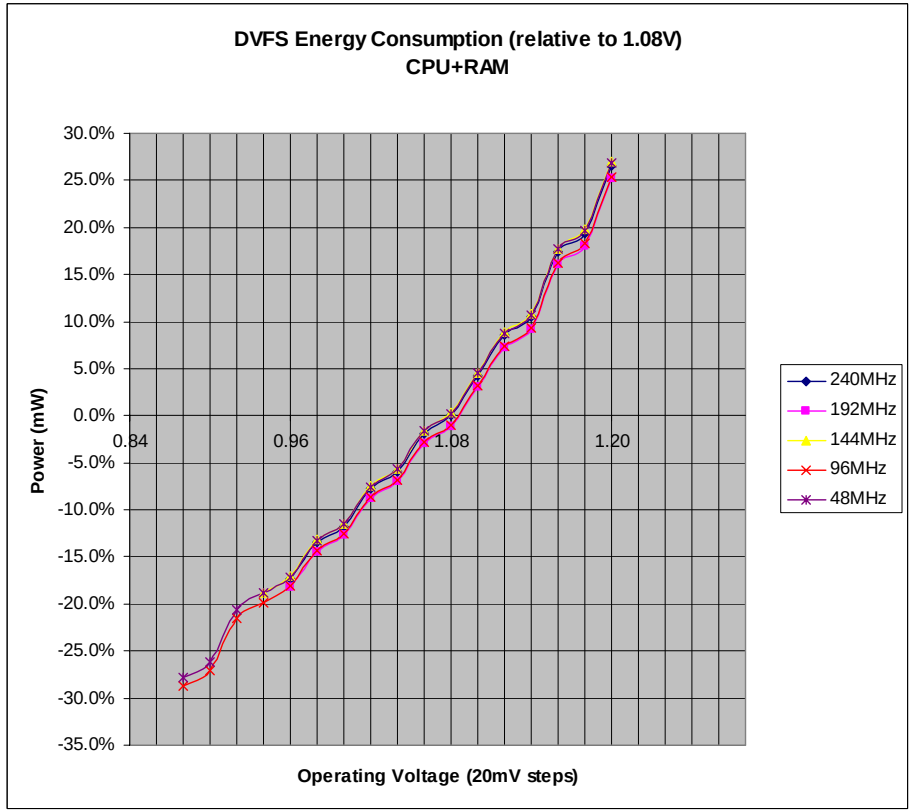

- 65LP process, typical silicon, 23 degC
- RAM not safely scalable below 1.08V
- Over 40% energy savings measurable from 1.2V

MP – Exploiting DVFS

MP extends control over power usage by implementing voltage and frequency scaling and turning off unused processors

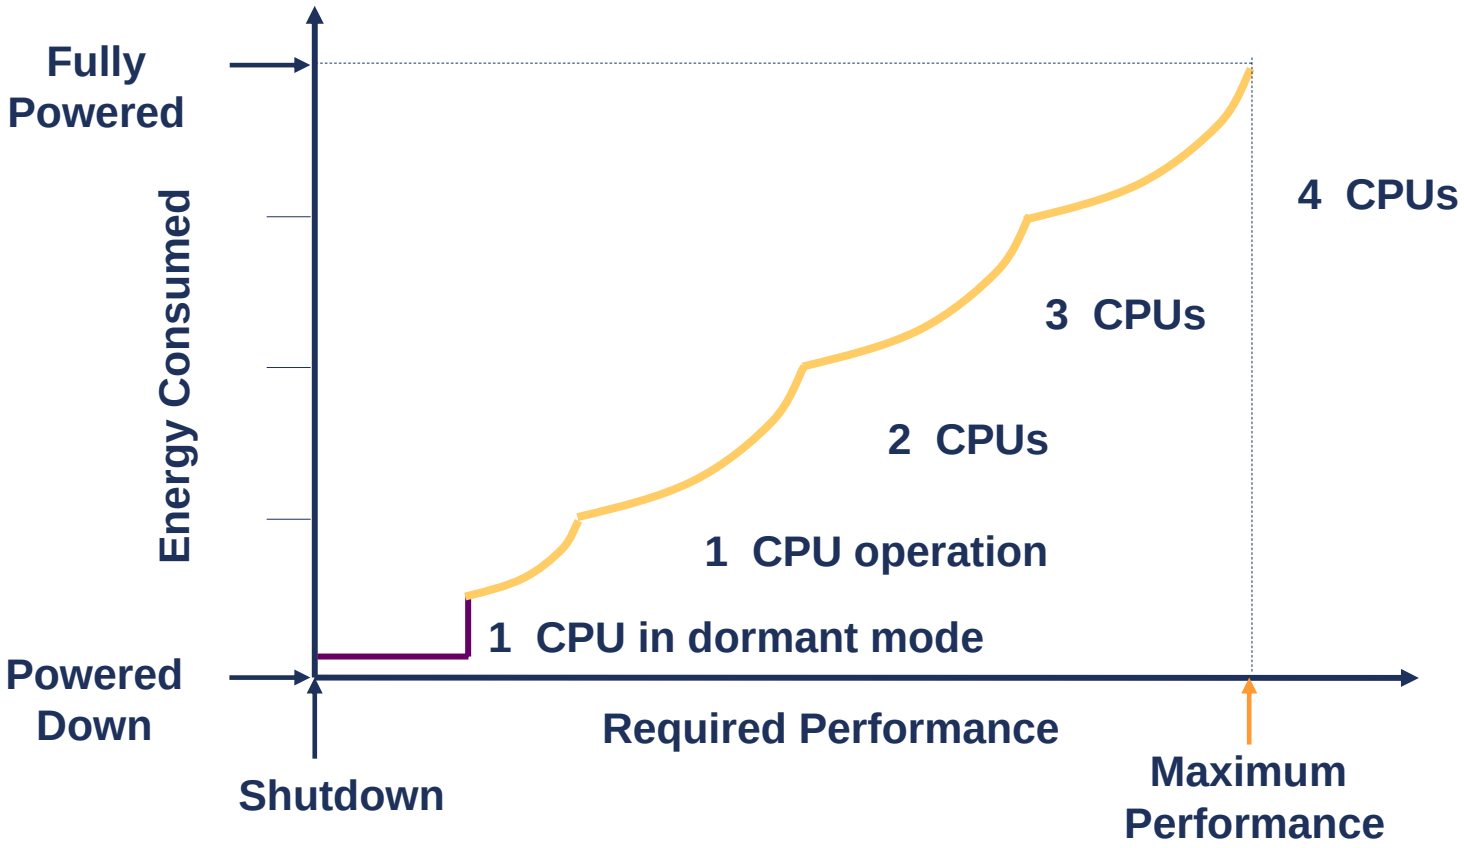

Workload Variance

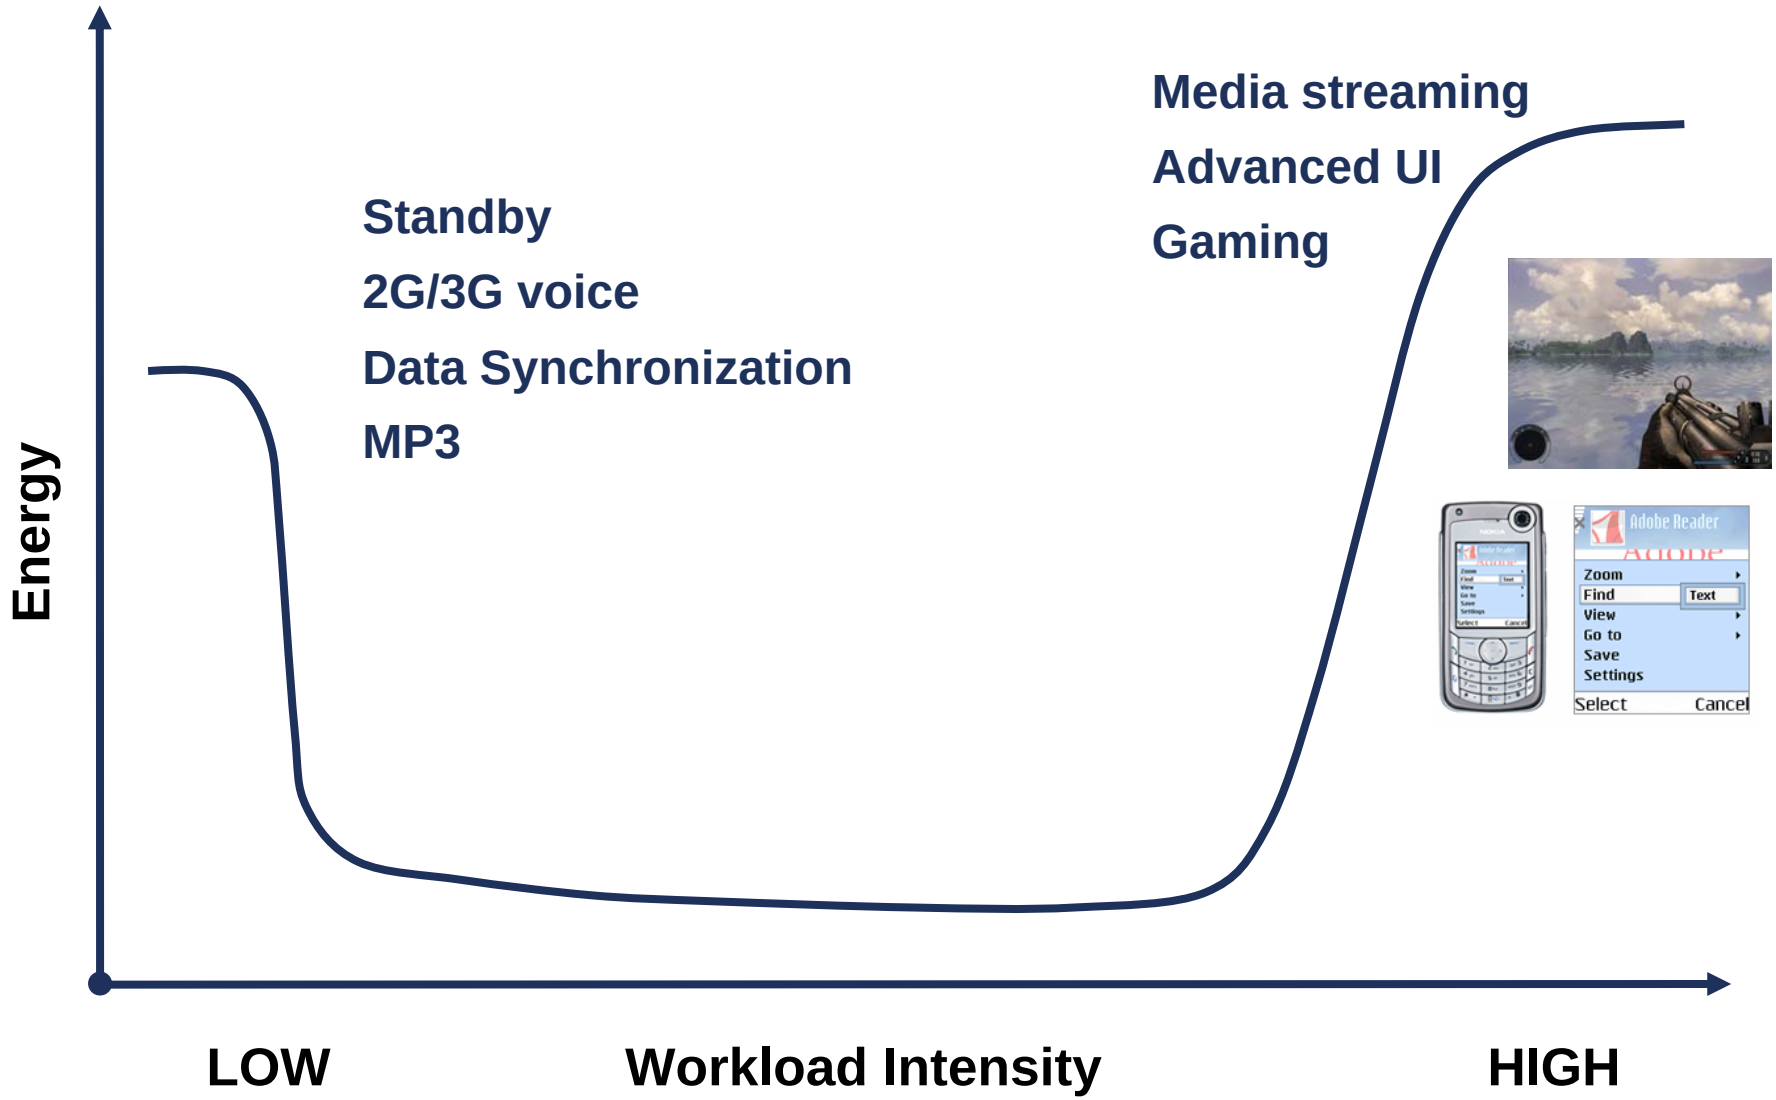

MP – Exploiting Workload Variance

- Enable tasks to migrate between cores with different performance capabilities and power footprints
- Minimise energy consumption using adaptive shutdown

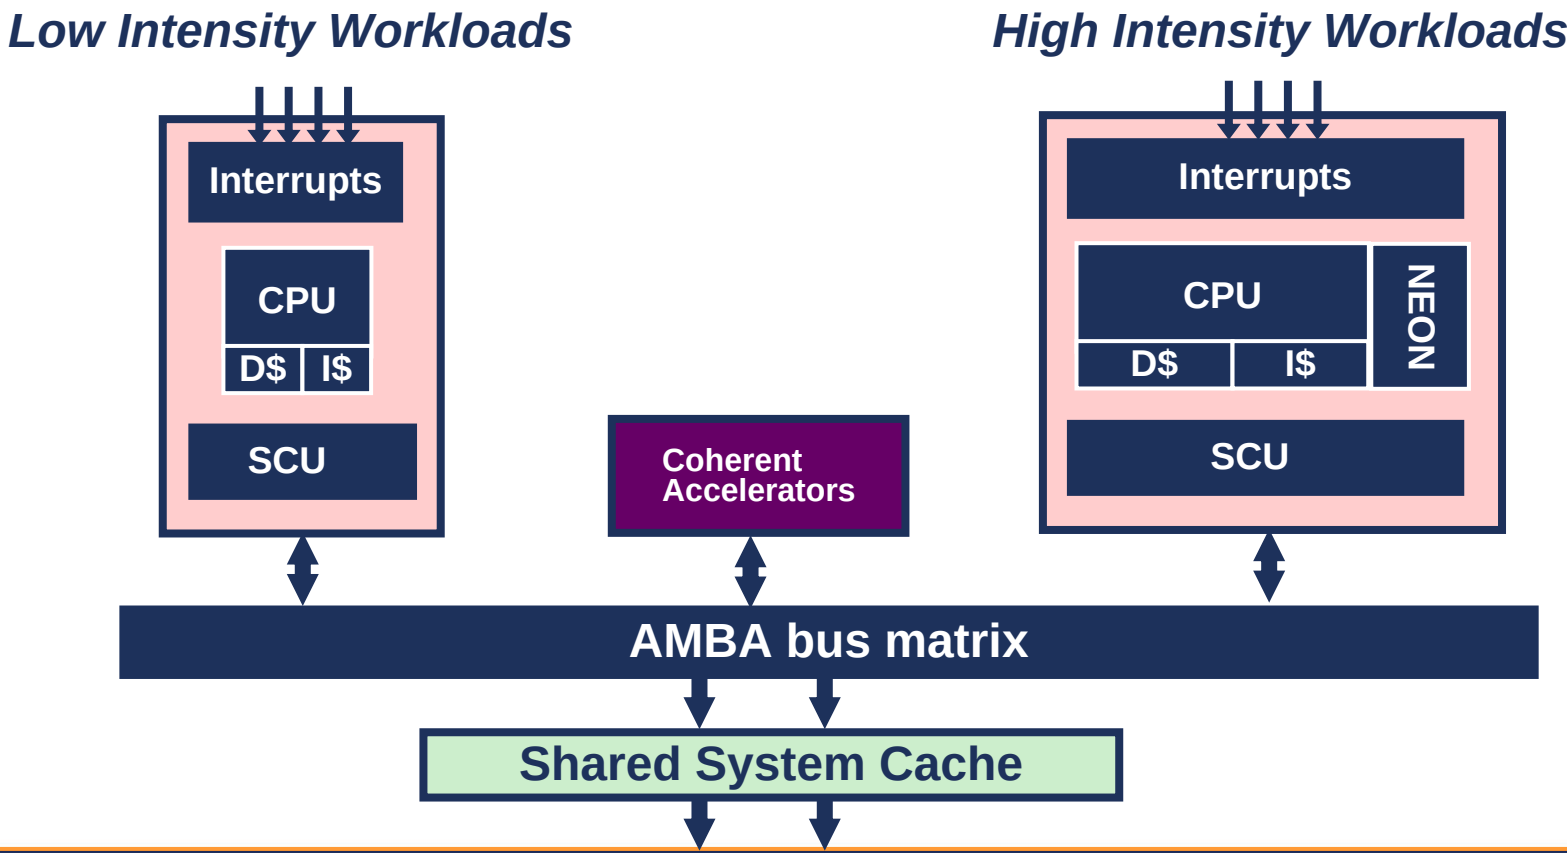

Exploiting Uncertainty

- Uncertainly leads to performance and power overheads
  - Increasing uncertainty with design scaling
  - Intra-die process/temperature variations, inductive noise
- Key Observation: worst-case conditions are highly improbable
  - Significant gain for circuits optimized for common case
  - Need mechanisms to tolerate infrequent worst-case scenarios

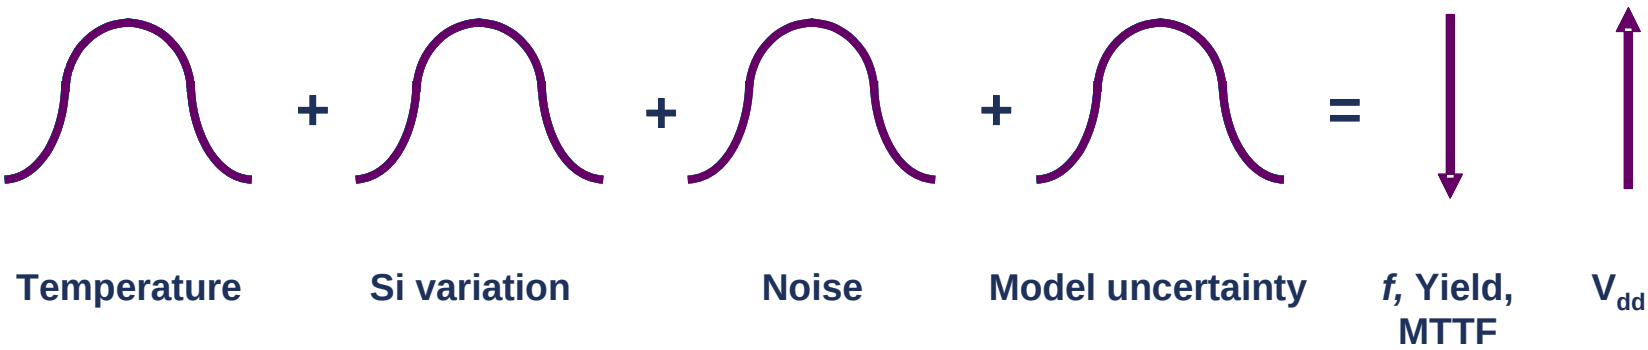

Razor – Beyond Worst Case Design

- Deep submicron processes are problematic
  - Process variation, IR-drop, temperature fluctuation, model uncertainty
- Standard solution: larger built-in safety margins
  - Yields lower performance, higher cost, higher power devices
- Most severe problems are also the least frequent
  - Soft errors, capacitive, inductive noise, charge sharing…
  - Only pay a penalty when the problem actually occurs!
- Break some conventions
  - Computation may not always be correct
  - Miss timing some of the time, compensate at run-time
- Speculate on correctness
  - Assume that circuits work as expected, recover if not
  - Speculation is key to minimizing run-time overhead

Energy Savings Using Razor

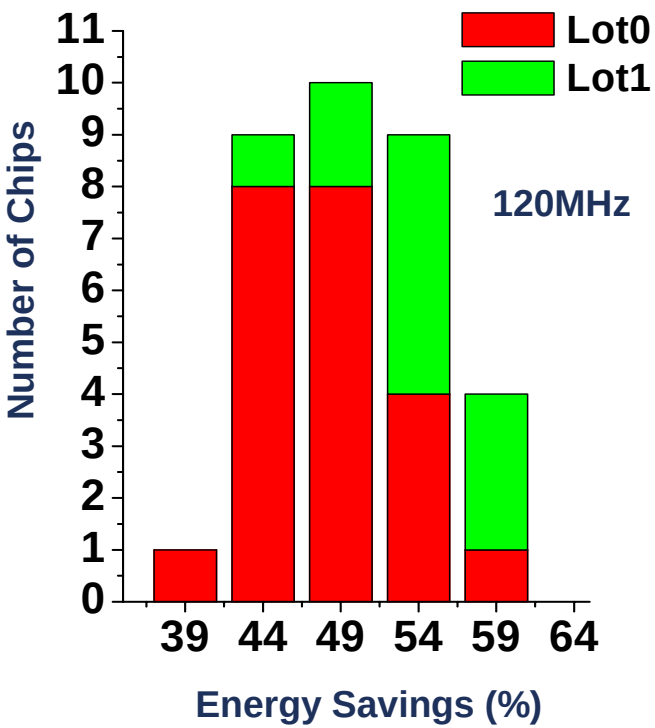

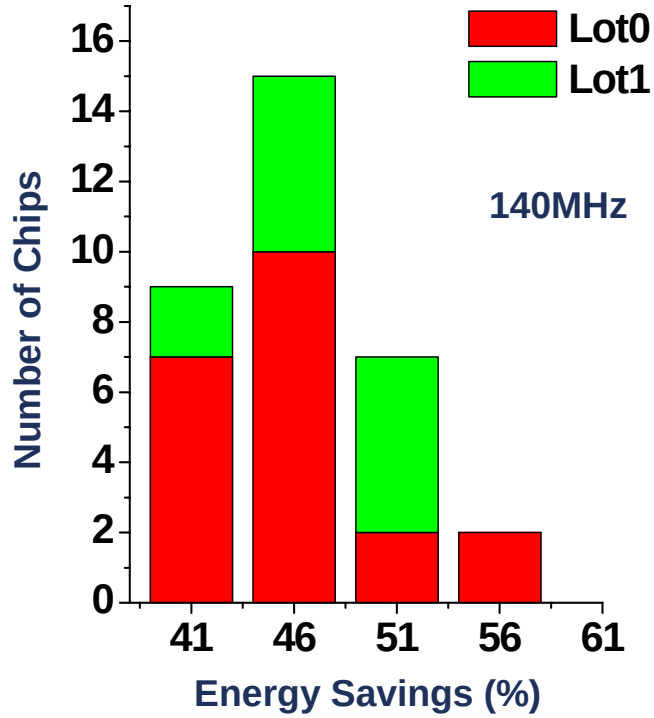

- Under typical case conditions all chips at least 39% more energy efficient
  - Worst-case design uses margins for corners that are very infrequent
- Typical-case operation requires an understanding of when and how systems break
  - Razor specifies the microarchitectural requirements

Summary

- Energy efficiency is now the limiting factor for many systems
- Industry challenge is to maximise the usable performance from a highly constrained energy envelope
- Market growth requires improved features and lower cost
- Heterogeneous MP clusters offer the potential for large increases in average energy efficiency
- Exploiting uncertainty can reduce energy through elimination of margins and correction of errors